

**Transcript of Integrated Review of the Capital Framework for Large Banks Conference:
Fireside Chat - Vice Chair for Supervision Michelle W. Bowman and Sam Altman, OpenAI
CEO
July 22, 2025**

VICE CHAIR FOR SUPERVISION BOWMAN. We're really looking forward to all of the conversations that we're having today, especially for what we're going to take away as we're continuing to do our work within the interagency and also with the Federal Reserve on capital and as we're looking to the future of banking more generally. I also want to take a moment to thank our panel participants before we get started with our fireside chat this afternoon. Thank you for taking time to be a part of this important conversation. As we're thinking about what regulation will look like in this space and in many others, we really look forward to engaging with you as we continue to have these discussions.

But today, in many ways, is about the future of banking. And with that in mind, we'd like to now turn to another influence that's shaping innovation and finance. While innovation has always played a role in the evolution of the banking industry, it's becoming clear that new technologies are not merely incremental improvements, but potentially huge leaps that could fundamentally alter the structure and function of our financial system.

One of these technologies, of course, is artificial intelligence. And I can't think of anyone who's better situated or prepared to discuss AI and the role of innovation in transforming finance and our economy more generally than Sam Altman, the chairman and CEO of Open AI. Sam, thank you for making time to be with us today. I want to welcome you. Thank you for being here.

SAM ALTMAN. Thank you very much for having me.

[Applause]

VICE CHAIR FOR SUPERVISION BOWMAN. It might be helpful just to sort of frame where we are with the landscape of AI and innovation more broadly. Would you just kind of set that framework for us?

SAM ALTMAN. Sure. Only five years ago, AI was still thought of as something that was in the distant future if it was going to happen at all. And even two and a half years ago, which was right around the time ChatGPT launched, it still hadn't moved past the sort of nerds in Silicon Valley. ChatGPT launched on November 30 of 2022. That was even before GPT-4. And since then, the progress has been quite rapid. The adoption, the economic impact has also started to be quite rapid. Just last week, we had a model that was able to achieve gold-level performance on the IMO. This is something that I think if you told most people in the field would happen, even a few years ago, they would say, absolutely not. That's as good as our smartest humans that are true experts in their field. We're now hearing from scientists saying they're two, three times more productive. We're hearing from computer programmers that say they're ten times more productive. It's completely changed what it means to write software. We have systems already that can perform at expert-level intelligence in many, many fields. Now, they cannot operate on tasks that are as long-horizon as humans can, so there's still a big limitation there. But even if progress were frozen right now, which of course it won't be, I think we still have years ahead of us for society and the economy to really digest this technology and figure out what the impact is going to be. There was a saying for a long time that I thought was great and we should try to get to again which is: "electricity too cheap to meter." And we didn't quite deliver on that as a society, although I think we still should. It does in fact look like we're about to deliver on intelligence too cheap to meter. We've been able to drive down the cost of each unit of

intelligence by more than a factor of ten each year for the last five years. Looks like we'll do that for the next five years too, maybe even more.

This weekend I used one of our upcoming models to do a computer programming task that I had wanted to do. I'm sort of like a home automation nerd, and I wanted the lights and music in my house to do this specific thing. I knew that before this technology it would have taken me days to do that. I was hopeful that with this technology, given our recent progress, I'd be able to do it in hours. I was able to do it in five minutes. The AI did almost all of the work. This is something that just a year ago you would have paid a very high-end programmer 20 hours, 40 hours, something like that to do. And an AI did it for probably less than a dollar's worth of compute tokens. So this is an amazing change. And the speed with which it has happened, the speed with which it will continue for the next few years, I think is still quite unappreciated. We weren't sure, even a year ago, how much further our current research roadmap was going to continue, if we're going to hit some sort of limit. At this point, it looks like we've got many years ahead of us of almost certain progress.

VICE CHAIR FOR SUPERVISION BOWMAN. Well, fantastic. That really helps us frame the next conversation, which is - -- You're talking to a room full of people in the financial industry and in banking, many of whom are already thinking about how they're going to use AI or already using AI. But how could you compare the potential for AI and productivity with other tech gains that we've seen in the past? I'm of the age that I was just starting business when the internet became something that we used more broadly. So how could you kind of frame that for us? And is there an analogy that you might use to describe where we are?

SAM ALTMAN. I've never seen a technology revolution quite like this. The historical examples are -- People talk about the industrial revolution; they talk about the computer

revolution. The Internet did change a lot of things, but I don't ever remember something that would have cost \$10,000 of knowledge work a year ago costing \$1 or \$0.10 or whatever it would cost now. For my programming example from a minute ago, I think this is just unprecedented.

Now it doesn't work like that for everything. Things in the physical world, robots, whatever, that's going to take a lot longer. Maybe in 2020 it cost \$10 to have an Uber drive you across town. Actually, let me say, maybe it cost \$100 to have an emergency parcel service delivered across some amount of time. And it cost \$100,000 to have an application written. In 2030, maybe to have that software application written costs literally \$0.10, \$100,000 down to \$0.10. And having that package delivered has gone from \$100 to \$1,000. And it's going to take a while for the robotics, for the full humanoid that can drive the car, take out the package, go upstairs, push the elevator button, the whole thing. That's just going to take a little while. But for tasks that can be done in front of a computer, it's quite -- I don't know any precedent for what's happening.

The analogy that I like the most, people say, you know, is this like the industrial revolution? Or, you know, is this like something else? The one I like is this is like the transistor. The transistor was a discovery of really deep science, physics, that turned out to be very hard to discover. Easy to understand once you understood it. Completely economically transformative. And the value sort of diffused throughout all of society as a massive productivity win. But there was a short period of time where there were a lot of transistors; they're called semiconductor companies, and it was this huge boom. And most of those are now kind of like in the background. You're all using devices with lots of transistors in them. They're all over this room. We don't think about this as a transistor device or that as a transistor device. You know, we think of this as like a microphone and a computer and a screen and a camera and whatever else. But

this technology was an incredible discovery. It changed what we were able to build. It became part of kind of everything pretty quickly. And you don't talk about companies much as transistor companies, except maybe companies like TSMC or [inaudible] or something like that. And in the same way, I don't think you'll be talking about AI companies for very long. You will just expect products and services to use this technology. You won't think about it that much. You will expect products and services you use to be smarter than you and that will be a normal part of the world.

Like the transistor, this is a technology that scales beautifully. So we had Moore's law with transistors. We have -- I don't think we have a name for it yet, but we have these scaling laws with AI where it just gets better and better and we learn how to really industrialize this and put this everywhere. So I think that's the best historical analog. And we did see phenomenal productivity gains because of the transistor. But again, what we're seeing now with AI, to say nothing of like the weird cases, like the fact that at some point AI can invent its own next versions and really accelerate progress even further. With what we're seeing right now, this just goes quite, quite far.

VICE CHAIR FOR SUPERVISION BOWMAN. Well, that's a great analogy. And I think when you said industrial revolution -- and one of the things that we think about here at the board is the importance of understanding how labor dynamics might change as a result of that. So could you give us your perspective on how you think the labor market and labor productivity or productivity more generally might be affected for some of the banks in the room or other industries that might be tuning in by our live stream?

SAM ALTMAN. I think that I like to remind people in our company of this, no one knows what happens next. There's a lot of these like really smart-sounding predictions. People make all kinds of like, oh, this is going to happen in this and the economy over here. We have no

idea. No one has that, in my opinion. No one really -- Like this is too complex of a system. This is too new and impactful of a technology. It's very hard to predict.

There are totally -- there are cases where entire classes of jobs will go away. There are entirely new classes of jobs that will come. And largely, I think this will look somewhat like most of history in that the tools people have to use their jobs will let them do more, achieve things in new ways. What it means to be a doctor or a lawyer or a computer programmer clearly will change, but people will still need medical care, and I want to talk to a human. People still need legal advice and they'll want someone that they can trust to be on their side. People will still want computers to do things for them. But the amount that one person can do and also the amount that we expect out of one person will be incredible. Like all throughout history, we have these new technologies come along and people say this is the end of work, this is the end of jobs. It turns out people seem to want unlimited stuff -

[Laughter]

SAM ALTMAN. And have a huge desire to express their creativity and be useful to people. And I'm still waiting for that promise from the industrial revolution that we only had to work four hours a week and got to play on the beach and hang out with our kids and here we all are.

But one thing I believe is you never fight biology. Evolution has had too long. We are too finely shaped. The sort of biological drives that we have and what it means to be human, technology doesn't change that. You can't beat that. And so I don't think the fundamental things that make us work and kind of animate society are going to go anywhere. And we'll probably all still be complaining that we have to work too hard even though we're unimaginably rich.

Probably if we could look from today at the people 100 years in the future, we would look at those jobs and say those are not real jobs. You are not actually busy. You have unimaginable luxury. You have everything you could possibly need. You have nothing to do, so you're making up a job to play a silly status game and to fill your time and to feel useful to other people. And those are exactly the things that someone from 100 or 500 years ago would say about us today. And that's just the way of things. Like, I think that's great.

[Laughter]

VICE CHAIR FOR SUPERVISION BOWMAN. That's a unique vision of the future. But if we dial it back to today, as regulators, we're always very risk-averse and we're also very cautious about the way that we're safeguarding our data and ensuring that when we can use innovation, as we're doing our work -- and I know banks are very interested in ensuring that they can leverage these kinds of technologies. But the data that we are protecting in our organizations is really critical that we're able to maintain its security. So how should we think about risks? And given that we're in government, in one section of government, how do you think about government uses?

SAM ALTMAN. So we expected that the financial sector and also the government itself were not going to be early adopters of our technology. AI, it's gotten much better, but when we first launched, the thing that the average person thought about AI is that it hallucinated a lot. I remember when we launched GPT-3, there was this whole survey of academics who were supposed to be experts in AI, and they were polled on what percent of GPT-3 answers were hallucinations. The real answer was 0.1% or maybe even less. But the consensus of the academic experts was 50%. They thought half the time people were talking about chat GPT, it was

completely hallucinating something. It was obviously not true, but that branding kind of stuck with us for a while. And so we assumed that financial services, to say nothing of the government, were not going to be our early adopters.

Some of our biggest early enterprise partners turned out to be financial institutions. Morgan Stanley, Bank of New York, these are major partners of ours that we're doing fantastic work with. And we were kind of like, are you all sure? And they were like, yeah. We really want to do this. And they have figured out how to use the technology, how to structure it enough that they can rely on it for critical -- Many other financial institutions too, rely on it for critical processes.

We are increasingly working with the government to roll out our services to lots of government employees, and a thing that someone said that stuck with me is, hey. We realize this is a new technology; we've got to put some new controls on it. But the risk of us not doing this -- If we don't adopt this, the risk is that we don't continue to exist as a business. We know we cannot compete with -- If we're a bank, we know we cannot compete with a new bank that is really going to build an AI-first experience and use AI for the entire stack of things. So it was one of the most innovative industries, and the adoption and rollout has worked better than I thought.

There clearly are risks to mitigate. We talked about hallucinations. A new one on the horizon is what's sort of called prompt injections, the idea that once a model really becomes personalized to you and your data, someone else can kind of trick it to tell someone something it shouldn't. I may know a lot of private information about you, but I understand when I can share that with him and when I can share that with her and that I should never share it with him. And the models -- that's a new thing that's happened. [Inaudible] the models get all of this personal

information in them. So there's some work to do there. But we've been able to kind of go through these risks carefully and manage them, but really deliver a lot of the upside.

VICE CHAIR FOR SUPERVISION BOWMAN. Speaking of personal information and those kinds of things, I might just pull a thread on that. And you may not know much about this, but this is something that the banking industry is really, really interested in right now is fraud and impersonating individuals to facilitate fraud. Are there ways that we could mitigate that kind of activity or that we should safeguard if there is the use of AI or identification of that type of impersonation?

SAM ALTMAN. Great, great question. I am very nervous about this. A thing that terrifies me is apparently there are still some financial institutions that will accept a voice print as authentication for you to move a lot of money or do something else. You say a challenge phrase and they just do it. That is a crazy thing to still be doing. AI has fully defeated that. AI has fully defeated most of the ways that people authenticate currently, other than passwords. But all of these fancy -- take a selfie and wave or do your voice or whatever, I am very nervous that we have a significant impending fraud crisis because of this. We have tried, I think other people in the industry have tried to warn people, hey. Just because we are not releasing the technology doesn't mean it doesn't exist. Some bad actor is going to release it. This is not a super difficult thing to do. This is coming very, very soon. There are obviously some reports now of these ransom attacks where people have the voice of your kid or your parent and they make this urgent call. That is going to get so compelling.

Society has to deal with this problem more generally, but people are going to have to change the way they interact. They are going to have to change the way they verify, like this person calling me. Right now it is a voice call. Soon it is going to be a video FaceTime. It will be

indistinguishable from reality. Teaching people how to authenticate in a world like that, how to think about the fraud; this is a huge deal.

VICE CHAIR FOR SUPERVISION BOWMAN. That might be something we could think about partnering on, the identification of those hallucinations or the intentional impersonations, I think would be a really beneficial thing to engage in. Being a parent of two teenagers, one just graduated from high school and the other is a sophomore in high school, apparently a lot of these kids these days are using chat GPT and other AI to facilitate their work and to complete high school. Do you have thoughts about that and how AI can be used in a beneficial way instead of in this way for kids and with teaching?

SAM ALTMAN. Let me tell two stories and then I will answer that. I never met my grandfather; he died before I was born. But a story that was told to me by my grandmother was about what happened -- he was very good at math -- what happened when the calculator came out and he always liked new technology. When the calculator came out, math teachers at the time apparently said, this is a disaster. This is the end of math education. If you don't have to learn how to use a slide rule or look up stuff in the logarithms books or whatever they did at the time, why bother teaching anyone math? These kids are never going to learn. It was apparently a real meltdown.

Of course what happened is with better tools, we use our brain power for something else. We start teaching calculus in high school. We start exposing people to more advanced math. My own version of the story was when I was in junior high and Google came out and the teachers at the high school, I heard this from the older kids, were really panicked because if there was this crazy new thing called Google, you didn't have to memorize facts. You didn't have to memorize the year that a certain war was fought. You could just look it up. And what was the point of

history class? If you didn't have to fire up your car and drive to the library and learn to use the card catalog and then walk down and try to find this book, but someone had checked it out, and then try to find some other book -- If you didn't have to suffer that hour, then what was the point of learning?

[Laughter]

I had gone to the library and used the card catalog as a little kid a little bit. I was like, no, that's a really bad experience. I could do something better with that hour of my life. And again, what happened -- So there was this brief effort at my school to ban Google and get these pledges not to use it and whatever, and then people realized, oh man. We can give our students more tools and we can expect more out of them. And sure, maybe they drive to the library a little bit less but they can use that time to think harder or come up with new ideas or do whatever, and that was really great. We expect more, we get more. The potential is higher but the expectations are much, much higher. And I think it's the same thing now with ChatGPT. It is totally true.

Back to that November 30 date when ChatGPT first came out, kind of the first mega passionate early adopters were for sure students using it to cheat on their finals or papers or whatever that December. And you had these school districts falling all over themselves to see who could ban ChatGPT the fastest. So we launched November 30. By December 7, some people have banned it. By December 14, it's really careening off the cliff. Schools leave for the break and it looks like ChatGPT is never going to be allowed inside of any educational institution again. It's a two-week-old product. It's done.

By the middle of January, you start seeing these things from principals, from superintendents, from serious, high-up people in the world of education saying, actually we made

a huge mistake. This is the best learning tool ever. Our students, at least our sort of self-starting students are using this to go learn anything. If we ban this from our schools, we're just not going to be globally competitive. And we've got to retool our curriculum because this is like the calculator, and now we have a calculator for words. So sure, maybe a take-home essay is no longer the right way to evaluate a student, but that process of learning to write an essay, that process of thinking better because you're like -- I still write things, but I never show anybody else, just to organize my thoughts. That's really important. So we said, great.

Now, the good news is students for sure have learned how to use this to learn better, to teach themselves to think things, to think in new ways. We're going to roll out a new way to help students learn better with ChatGPT very soon.

The bad news is on the whole, in those two years, two and a half years since, the curriculums have not moved as much as we might have hoped. People seemed to have a lot of energy that they were going to have to start educating and evaluating in new ways. And then kind of the molasses of the education system did its thing and we're still giving a lot of take-home essays. And I think that's a losing battle. We should be teaching people to use -- like give them assignments that require them to use tools like ChatGPT. Cannot be done without it, that's fine, but expect much more out of them. And I still hope we're going to get there.

VICE CHAIR FOR SUPERVISION BOWMAN. Sounds like a transformation of our education system.

SAM ALTMAN. But it has to be. These kids are going to grow up and be adults with unbelievably powerful AI. And if we're not training them for that world, we've really missed the mark.

VICE CHAIR FOR SUPERVISION BOWMAN. That's very helpful. Thank you. I'll make sure my daughter understands what our expectations are as she's finishing high school.

But to move back to some things that you've talked about in the past and to move a little bit back to the business community, you've said that AI will make small businesses -- Or it's going to create an environment for entrepreneurs that's going to be second to none. A new era, I think, is what you called it. How do you see small businesses? That's something that a number of our board members are really engaged in and really, we care about making sure that there's new business formation and that there's opportunity there for new entrants in that. How do you see that being transformed? And not just for the Silicon Valley, small businesses, or new entrepreneurs there, but more broadly.

SAM ALTMAN. So the first thing we saw with ChatGPT, as I mentioned, was students. One of the coolest things we saw pretty quickly after was -- Everybody who worked at OpenAI started to get these stories of how people were using ChatGPT when it was still very novel in their business life. And everybody has their favorite, but I'll tell you my favorite.

I was in an Uber, and the Uber driver was telling me about this amazing new thing called ChatGPT and asking if I'd heard of it. And I was like, oh yeah. What do you do with it? And he was like, it's the craziest thing. I have this small business. It wasn't quite working, but now I have an employee in every area. It does contracts for me. It answers my customer support emails. It helps me come up with my marketing stuff. It designs my - He goes down this list. And he basically had this business running on ChatGPT. This was before it was a common thing. This was before GPT -- This was an extremely early adopter. But he had figured out how to run a business on ChatGPT. And it was not like he was taking jobs from other people. His business just would have failed. He couldn't pay for the lawyers. He couldn't pay for the customer support

people. He didn't know how to get someone to design the advertisements for him. He didn't know how to have a system that could automatically bid for ads on the internet for him. And ChatGPT just did all of this. Now obviously it's much better now. Now we have a whole industry around us that builds tools with our API to make people -- really truly one-click, automate all those things I just talked about. But the creative spirit of people that were just doing this inside of ChatGPT, like way back in the Stone Age, really stuck with me. And now it feels like it's on easy mode and people are doing that in all sorts of amazing ways.

VICE CHAIR FOR SUPERVISION BOWMAN. It's really transformational, I think, for me. The first time I used it, I was trying to write a haiku for a Kanpai toast at an innovation conference, which actually didn't go very well.

[Laughter]

VICE CHAIR FOR SUPERVISION BOWMAN. But at the urging of my family, of my teenagers and my husband who's involved in AI too, we tried to make it work and then went back to the drawing board.

SAM ALTMAN. The new version should be much better haikus. It's very good at creative writing.

VICE CHAIR FOR SUPERVISION BOWMAN. I'll think about that. But I wanted to make sure we have some time for audience questions if you're up for that. I see Anil in the front row. Anil Kashyap with Chicago Booth.

ANIL KASHYAP. Hi. So I assume one of the things that financial institutions will begin doing is using this to mine their data, maybe to do credit scoring, credit evaluation. What do you

say to the people that are worried that it's going to be picking up patterns and making decisions on things we wouldn't want to? And how are you guys thinking about protecting against that?

SAM ALTMAN. This is deeply outside my area of expertise, so I'll answer in the general case. One of the things that is cool about these models is they understand us in natural language and they're so steerable. So if you say, hey. You can go look at all this data to make a decision, but don't consider X, Y, or Z at all, like don't let that have any factor, generally speaking, it'll really be good at following that. Doing this in language models is different from the way that we used to do this sort of like just group all this data together into these clusters. And so you can really instruct and it will really, really follow your intent.

The second thing is, I think the question behind the question there is that these models are going to be very biased in ways we don't like. I think humans are quite biased. I think AIs are dispassionate and unemotional and don't have the kind of like built up, you know, whatever that people do. And I think it'll be possible for AI, correctly built, to be a significant debiasing force in many industries. And I think that's not what many people thought, including myself, with the way we used to do AI, but on the current trajectory, that seems very possible.

STEPHEN SCOTT. Hi, Sam. Stephen Scott with a company called Starling. I recently had the opportunity to interview Craig Mundie, who, you know, wrote a book with Henry Kissinger and Eric Schmidt, where they describe AI as almost the evolution of a new species. And Craig was saying that we can teach AI human moral codes. They can learn it and reflect it. So there's a lot of talk about the ethics of AI. My question for you is, will AI evolve its own morality?

SAM ALTMAN. I think people get very confused about whether they want AI to be a tool or a creature. And I don't, like, I'm very much in the tool camp. I don't think AI will have independent morality. I think that AI can certainly learn, you know, really study humanity and the best of our thinking. It can probably help us point out, here's a problem in your thinking. Here's a thing that should be different; here is like a real moral gap. And I don't know what the giant moral gaps are in our current worldview, but I guarantee you there are some. And if it can help us find those a little bit faster, I think that would be -- If it could point out those inconsistencies, I think that would be helpful. But I think it is easy to sort of impute creature-ness on AI that is not there.

WILSON ERVIN. Thanks. I read recently -- It's weird to use the phrase "traditional internet," but there are some people who say that AI will sort of destroy or really reshape the traditional internet, whether it's search, whether it's lots of other things. And obviously there's a huge economy now that we built on that. What's your sense of the amount of change that's going to happen in that world where AI will be most disruptive and where things will be more status quo?

SAM ALTMAN. I do think there are all of these things -- So I think AI will be somewhat disruptive to the way people currently use technology. The reason I was smiling is there's this funny thing that happens. Older people or people who are used to like emailing with a certain kind of etiquette will type out the bullet points that they want to communicate, put it into ChatGPT, ChatGPT will write this long, nice formal email with a bunch of like fluff and the bullet points somewhere in there, they'll send it across to somebody else, that person will put it back into ChatGPT and say please summarize this for me.

[Laughter]

SAM ALTMAN. And on the whole, the high school students think this is ridiculous and just say, just send the bullet points. That kind of like formal email thing is dead, just like -- You know, it's you're generating on one side; you're collapsing it back. This is ridiculous.

And it's funny, but I think there's something about the way we use the whole internet that is encapsulated there. When I wake up in the morning, I go through a bunch of apps. I read messages across five or six different things. I go check, you know, a thing here, a thing there, a thing there. And what I really would like -- And then my phone, all day long, it's just like blowing up at me and it feels like I'm like walking down the Las Vegas strip and these things flashing at me and it's very distracting, people like shoving things in my face.

And what I would like is my AI agent to be off using the internet for me, knowing when to interrupt me. It can tell when I'm focused working. It can tell when I'm in a meeting. It can tell when I got some time to think. It can override if necessary. Otherwise it can like nicely summarize stuff, respond to things for me, pull the things together. And, you know, I want, like, that bullet point, condensed thing. And I don't want the fluff, and I don't have to like go around and click around, and I don't have to, you know, respond to stuff I don't want to respond to. But that change is probably fairly disruptive to the way that the internet works now. And I think there'll have to be new business models that go along with that.

Like, you'll have to have new ways of paying for content. I have always wanted micro payments for content on the internet. I hope that finally happens. Maybe there'll be like new ways that we actually reduce, you know, spam and message overload with new kinds of protocols. But it does seem like we're heading towards a very different way. You will start your morning with technology when you wake up.

PETER HOOPER. I'm Peter Hooper, Deutsche Bank. Sam thank you very much for prolonging a career, at least raising productivity as the years advance rather than seeing a drop-off. You mentioned you're expecting significant job losses and significant job gains. Could you talk a little bit more about the areas and the potential disruption that could cause?

SAM ALTMAN. One thing I believe, just as a general statement first, is that we have no idea really, kind of like how much more labor supply it would really take to meet true demand today.

You know, when you are sitting in a doctor's office waiting room for an hour, I think that just means under supply of doctors, or that the doctors aren't productive enough. And, you know, it would be great if the doctor's waiting around, and as soon as you get there, like they're ready to see you and they have the whole thing ready to go. Lots of other examples like this. Every time you're, you know, wasting your time in any way, every time you're kind of clicking around the internet, can't quite do the productive thing -- I think we are in an under supply of labor to a degree that is going to look horrible in retrospect.

Now some areas, again, I think just like totally, totally gone. I don't know if any of you have used one of these, like, AI-customer support bots, but it's incredible. A couple years ago you'd like call customer support. You, like, go through a phone tree, you talk to four different people, they do the thing wrong, you call back again, you wait through it. It's like hours of pain, ton of time on hold, and the thing that you want doesn't happen. Very frustrating experience.

Now you call one of these things and AI answers. It's like a super smart, capable person. There's no phone tree; there's no transfers. It can do everything that any customer support agent at that company could do. It does not make mistakes. It's very quick. You call once; the thing just

happens; it's done. Answers right away, great. Now I don't want to go back there and I also don't -- it doesn't bother me, at all, that that's an AI and not a real person. So that's a category where I would just say, you know what? When you call customer support, you're going to be talking to an AI, and that's fine.

A lot of these other things -- I really do want a human doctor. ChatGPT today, by the way, most of the time can give you better -- it's like a better diagnostician than most doctors in the world. And I'll, like many people here, probably put my symptoms in and test results, and you know, like -- There's all these stories on the internet of ChatGPT saved my life, and you know, I had this rare disease and it found it and all these doctors didn't do it, and yet people still go to doctors. And I am not like -- Maybe I'm a dinosaur here, but I really do not want to like entrust my medical fate to ChatGPT with no human doctor in the loop. Would anybody here rather just have ChatGPT diagnose them than a doctor? Even though you know it's better? That's like quite interesting, right? So this is a class where we're going to keep doing things largely the way that we did.

We talked earlier about this example of computer programmers. Like again I think it's amazing that a computer programmer is now ten times more productive. Salaries of programmers going up extremely rapidly in Silicon Valley, so are expectations. It turns out, I think, the world wants a gigantic amount more software, hundred times, maybe a thousand times more software. So maybe each person can now write ten times as much software. They're going to make three times as much. The world will be happy because the world is running way more software. The programmers will be happy, too, and I think we'll see many categories like that. Things in the physical world will keep being done by humans for a while, but when this robotics wave comes

crashing in another three to seven years, I think that's going to be a really big thing for society to reckon with. Yeah, that's kind of how I summarize it.

VICE CHAIR FOR SUPERVISION BOWMAN. I think we have two more questions over here and we might have time for a third, but let's go ahead. Rob, you already have the microphone, would you introduce yourself, please?

ROB BLACKWELL. Sure. Rob Blackwell with IntraFi. We've been raised on decades of sci-fi telling us that AI is eventually going to kill us all. And since you know more about AI than arguably anybody in this room, I just want to ask you, what does keep you up at night? What are the things that you worry about when it comes to AI, and how do we prevent those things that you worry about from coming true?

SAM ALTMAN. I think there's three sort of scary categories. There's a bad guy gets superintelligence first and misuses it before the rest of the world has a powerful enough version to defend. So, an adversary of the U.S. says I'm going to use this superintelligence to design a bioweapon to take down the United States power grid, to, you know, break into the financial system, and take everyone's money. Something that would just be hard to imagine without significantly superhuman intelligence, but with it becomes very possible, and because we don't have that we can't defend against it. So that's category -- broad category one. And I think that the bio capability of these models, the cybersecurity capability of these models, these are getting quite significant. You know, we continue to like flash the warning lights on this. I think the world is not taking us seriously. I don't know what else we can do there, but it's like -- This is a very big thing coming.

Category two is that sort of broadly called loss of control incidents, where the -- That's kind of like the sci-fi movie. The AI is like, oh. I don't actually want you to turn me off. I'm afraid I can't do that, you know, whatever. And that's -- I think that is less of a concern to me than the first category, but a very grave concern if it came to pass. There's a lot of work we and other companies do on model alignment to prevent that from happening, but as these systems become so powerful that's a real concern.

And then there's the third one, which I think -- those first two are sort of easy to think about and imagine. The third one is, to me difficult, more difficult to imagine, but quite scary, and I'll explain what it is and then I'll give a short-term and a long-term example. This is the category where the models kind of accidentally take over the world. They never wake up; they never do the sci-fi thing; they never open the pod bay doors. But they just become so ingrained in society and they're so much smarter than we are, and we can't really understand what they're doing, but we do kind of have to rely on them. And even without a drop of malevolence from anyone, society can just veer in a sort of strange direction.

When I was a kid and Deep Blue, that AI system built by IBM, beat Garry Kasparov in chess, I remember my dad saying, this is the end of chess and no one's going to play it again. But then it turned out that actually although the AI was better than humans, AI plus a human together was way better than an AI or the human. You know the AI would present ten options and the human would pick the best one or something like that and play the move. And everybody said, oh. We have this wonderful future of man and machine together. It's all -- no problem, whatever. That lasted two months, three months, something like that, and then the AI got so smart that the human only made it worse because they didn't understand what was really going on, and the AI alone trounced the AI and human. It's been like that ever since.

Now another interesting part of that story is everybody was convinced in the '90s that was the end of chess because if AI could beat humans, why should humans care? Chess has never been more popular than it is today. People love to watch chess. We're very focused on real people doing real-people things, so there was a very interesting thing that happened there.

But this phenomenon I think is a really big deal. In the short term you can see it where people may be -- We call this emotional over-reliance. People rely on ChatGPT too much. There's young people who just say like, I can't make any decision in my life without telling ChatGPT everything that's going on. It knows me; it knows my friends. I'm going to do whatever it says. That feels really bad to me, and is a really common thing with young people, and we're studying that. We're trying to understand what to do about it. Even if ChatGPT gives great advice, even if ChatGPT gives way better advice than any human therapist, something about kind of collectively deciding we're going to live our lives the way that the AI tells us feels bad and dangerous and a bunch of things like that.

The longer-term category is, you know -- back to that chess example -- what if AI gets so smart that the President of the United States cannot do better than following ChatGPT-7's recommendation but can't really understand it either? What if I cannot make any better decision about how to run OpenAI and I just say, you know what? I fully hand it over. ChatGPT-7, you are in charge. Good luck. That might be the right decision in any individual case, but it means that society has, like, collectively transitioned a significant part of decision-making to this very powerful system that is learning from us, improving for us, evolving with us but in ways we don't totally understand. So that's the third category of how I think things can go wrong.

VICE CHAIR FOR SUPERVISION BOWMAN. I think there was one question behind you, Rob, and then I might just ask you to wrap up by helping us encourage innovation in the banking system and how you might recommend we do that.

JOE CAVATONI. Hi, Joe Cavatoni with the World Gold Council. I think you're kind of touching on it in your third example that you've just given. Can you give us a little bit of thinking on what you're thinking about developed markets and the benefits that come to develop markets with AI, also developing markets, because this is potentially a level-setting phenomenon that could develop, correct?

SAM ALTMAN. Very much. I think it'll be a deeply level-setting phenomenon. Like, you know, maybe you don't love the healthcare or the banking system in the United States, but at least you have one. You know, at least you can get financial advice. At least you can get medical advice.

In a lot of the developing world, the alternative to a ChatGPT doctor is not a real doctor, it's nothing at all, and then you'd definitely rather have this. And if this keeps getting better fast enough, maybe you just get better service in the developing world in a bunch of ways. But I am very interested in what it means to give everybody on Earth like a free copy of GPT-5 running for them all the time. Every business really enabled with this level of technology to be able to give better financial advice, to detect fraud better, to underwrite risk better. Again, watching what is possible now makes me very optimistic.

In the developed world, I think the biggest challenge will be risk tolerance and regulation, for very good reasons -- but about how quickly do you want to adopt these things. And I think like we've seen with a few other technologies, in much of the developing world, people will just

skip a few generations. They'll feel like, you know, mobile and the internet sort of -- and they'll go right to just, you know, we're going to run everything on AI and we're going to deliver goods and services, at least services, at 1/100 of the cost. And I think you'll see some economies transform there very quickly.

I do think that in -- You know, obviously I'm most familiar, by far, with how things work in the United States, but I do think financial services here should just be completely different in another decade. Like the way that we move money, the way that we give financial advice, the way that we think about underwriting risk, it just seems like we can so dramatically improve all of that so quickly now.

VICE CHAIR FOR SUPERVISION BOWMAN. How do you see open AI or your services facilitating that evolution? And how can we as regulators think about what the constructs should be as we're implementing or allowing for innovation to occur?

SAM ALTMAN. I promise I didn't come here for the TV commercial, but we'd love to work with any of you, of course.

[Laughter]

SAM ALTMAN. Whether it is us or one of our competitors, I think that -- Just in the last six months, as reasoning models have really rolled out this transition from these models that could think not at all and had to give an instant response to models that can think for many seconds or many minutes, and the robustness and reliability that has come with that, means that for an industry like this, I think the technology is finally really usable.

And a lot of people have not tried the latest generation of models. But I think if you do, you'll be like, oh. This is much smarter than most people.

And then on the government side, I think it's the same thing. Like, government has got to embrace this technology and will be able to do everything better.

VICE CHAIR FOR SUPERVISION BOWMAN. Well, thank you, Sam, so much for joining us. We ran a little bit over time and I'm sure everybody's ready for lunch. Thank you very much. Sorry for keeping you.

[Applause]